

This listing of claims will replace all prior versions, and listings, of claims in the application.

Listing of Claims:

1. (Previously Presented) A method implemented in a communication switch comprising:
 - accepting a packet associated with a flow;
 - generating a flow group identifier from the flow;
 - determining whether other packets associated with the flow group are present in a switch fabric;
 - assigning the packet to a path being used by the flow group if other packets associated with the flow group are present in the switch fabric; and
 - assigning the packet to a path using path congestion status information if other packets associated with the flow group are not present in the switch fabric.
2. (Previously Presented) The method of claim 1 wherein generating a flow group identifier from the flow comprises hashing a flow identifier.
3. (Previously Presented) The method of claim 1 wherein determining whether any other packets associated with the flow group are present in a switch fabric comprises maintaining an outstanding packet counter.
4. (Previously Presented) The method of claim 3 wherein the outstanding packet counter is associated with the flow group identifier.
5. (Previously Presented) The method of claim 4 wherein maintaining an outstanding packet counter comprises incrementing the outstanding packet counter each time a packet belonging to the flow group is sent into the switch fabric, and decrementing the outstanding packet counter each time a packet belonging to the flow group leaves the switch fabric.

6. (Previously Presented) The method of claim 5 wherein decrementing the outstanding packet counter is performed in response to a message from an output port.

7. (Previously Presented) The method of claim 6 further comprising:
transmitting the message from the output port to a corresponding input element,
transmitting the message from the corresponding input element, through the switch fabric, to an output element, and
transmitting the message from the output element to another input element corresponding to the output element, wherein the other input element originated the packet.

8. (Previously Presented) The method of claim 3 wherein maintaining the outstanding packet counter comprises resetting the outstanding packet counter if it remains non-zero for more than a predetermined period of time.

9. (Previously Presented) The method of claim 1 wherein assigning the packet to a path using path using path congestion status information comprises:
selecting a switch plane having at least one uncongested path, and
selecting an uncongested path of the selected switch plane.

10. (Previously Presented) The method of claim 9 wherein selecting a switch plane having at least one uncongested path is performed using a round robin discipline.

11. (Previously Presented) The method of claim 9 wherein selecting an uncongested path of the selected switch plane is performed using a round robin discipline.

12. (Previously Presented) A method implemented in a communication switch comprising:
accepting a packet associated with a flow;
generating a flow group identifier from the flow;

storing a first data structure in a switch fabric, the first data structure comprising a first entry the first entry comprising the flow group identifier, an outstanding packet indicator, and a path identifier; and

storing a second data structure in the switch fabric, the second data structure comprising a second entry, the second entry comprising the path identifier and path status information, wherein the path status information comprises an indication of whether a path failed and an indication of whether the path is congested.

13.-14. (Cancelled)

15. (Previously Presented) An apparatus comprising:
means for accepting a packet associated with a flow;
means for generating a flow group identifier from the flow;
means for determining whether other packets associated with the flow group are present in a switch fabric;

means for assigning the packet to a path being used by the flow group if other packets associated with the flow group are present in the switch fabric; and

means for assigning the packet to a path using path congestion status information if other packets associated with the flow group are not present in the switch fabric.

16. (Previously Presented) The apparatus of claim 15 wherein the means for generating a flow group identifier from the flow comprises means for hashing a flow identifier.

17. (Previously Presented) The apparatus of claim 15 wherein the means for determining whether any other packets associated with the flow group are present in a switch fabric comprises means for updating an outstanding packet counter.

18. (Previously Presented) The apparatus of claim 17 wherein the outstanding packet counter is associated with the flow group identifier.

19. (Previously Presented) The apparatus of claim 18 wherein the means for updating an outstanding packet counter comprises means for incrementing the outstanding packet counter each time a packet belonging to the flow group is sent into the switch fabric, and decrementing the outstanding packet counter each time a packet belonging to the flow group leaves the switch fabric.

20. (Previously Presented) The apparatus of claim 19 wherein the outstanding packet counter is decremented in response to a message from an output port.

21. (Previously Presented) The apparatus of claim 20 further comprising:
means for transmitting the message from the output port to a corresponding input element,
means for transmitting the message from the corresponding input element, through the switch fabric, to an output element, and
means for transmitting the message from the output element to another input element corresponding to the output element, wherein the other input element originated the packet.

22. (Previously Presented) The apparatus of claim 17 wherein the outstanding packet counter is reset if the outstanding packet counter remains non-zero for more than a predetermined period of time.

23. (Previously Presented) The apparatus of claim 15 wherein the means for assigning the packet to a path using path congestion status information comprises:
means for selecting a switch plane having at least one uncongested path, and
means for selecting an uncongested path of the selected switch plane.

24. (Previously Presented) The apparatus of claim 23 wherein the means for selecting a switch plane having at least one uncongested path uses a round robin discipline.

25. (Previously Presented) The apparatus of claim 24 wherein the means for selecting an uncongested path of the selected switch plane uses a round robin discipline.

26. (Previously Presented) A method for alleviating head-of-line blocking implemented in an input-buffered switch, wherein the switch comprises an input module, the input module including virtual output queues and virtual path queues, the method comprising:

assigning an incoming cell to one of the virtual output queues using cell destination information;

providing a head-of-line cell of the one of the virtual output queues to one of the virtual path queues using a dynamic hashing scheme;

selecting, for a switch plane link, one of a number of virtual path queues having a cell; and

sending the cell from the selected one of the number of virtual path queues over the switch plane link.

27. (Cancelled)

28. (Previously Presented) The method of claim 26 further comprising:

determining whether the cell sent over the link is the last cell of a packet; and

sending another cell associated with another packet to one of the virtual path queues if the cell sent over the link is the last cell of the packet.

29. (Previously Presented) An input module implemented in a switch, the input module comprising:

a plurality of virtual output queues for accepting cells based on cell destination information; and

a plurality of virtual path queues for accepting head-of-line cells from the plurality of virtual output queues, wherein each of the virtual path queues accepts a head-of line cell from a virtual output queue based on a dynamic hashing scheme.

30. (Previously Presented) The input module of claim 29 wherein the number of the plurality of virtual output queues equals a number of output ports of the switch.

31. (Previously Presented) The input module of claim 29 wherein the number of the plurality of virtual path queues equals a number of paths through a switch fabric of the switch.

32. (Previously Presented) The input module of claim 29 wherein the number of the plurality of virtual path queues equals the number of switch planes of a switch fabric of the switch multiplied by the number of paths through each of the switch planes.

33. (Previously Presented) A communication switch for multicasting a packet, the communication switch comprising:

- a switch fabric; and

- a traffic manager module configured to:

- generate a flow group identifier from a flow;

- determine whether other packets associated with the flow group are present in the switch fabric;

- assign the packet to a path being used by the flow group if other packets associated with the flow group are present in the switch fabric; and

- assign the packet to a path using path congestion status information if other packets associated with the flow group are not present in the switch fabric.

34. (Previously Presented) The communication switch of claim 33 wherein the traffic manager module is configured to generate a flow group identifier from the flow by using a hash of a flow identifier.

35. (Previously Presented) The communication switch of claim 33 wherein the traffic manager module is configured to determine whether any other packets associated with the flow group are present in a switch fabric by using an outstanding packet counter.

36. (Previously Presented) The communication switch of claim 35 wherein the outstanding packet counter is associated with the flow group identifier.

37. (Previously Presented) The communication switch of claim 36 wherein the traffic manager module is configured to increment the outstanding packet counter each time a packet belonging to the flow group is sent into the switch fabric, and decrement the outstanding packet counter each time a packet belonging to the flow group leaves the switch fabric.

38. (Previously Presented) The communication switch of claim 37 wherein the outstanding packet counter is decremented in response to a message from an output port.

39. (Previously Presented) The communication switch of claim 38 wherein the traffic manager module is further configured to:

transmit the message from the output port to a corresponding input element,
transmit the message from the corresponding input element, through the switch fabric, to an output element, and
transmit the message from the output element to another input element corresponding to the output element, wherein the other input element originated the packet.

40. (Previously Presented) The communication switch of claim 35 wherein the outstanding packet counter is reset if the outstanding packet counter remains non-zero for more than a predetermined period of time.

41. (Previously Presented) The communication switch of claim 33 wherein the traffic manager module is configured to assign the packet to a path using path congestion status information by:

selecting a switch plane having at least one uncongested path, and
selecting an uncongested path of the selected switch plane.

42. (Previously Presented) The communication switch of claim 41 wherein the traffic manager module is configured to use a round robin discipline to select at least one of the switch plane and the uncongested path.